JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1979, 31, 351-371 NUMBER 3 (MAY)

PARTIAL AVOIDANCE CONTINGENCIES: ABSOLUTE OMISSION AND PUNISHMENT PROBABILITIES[1]

BARBABA L. FLYE[2] AND JOHN GIBBON

COLUMBIA UNIVERSITY AND NEW YORK STATE PSYCHIATRIC INSTITUTE

Avoidance contingencies were defined by the absolute probability of the conjunction of responding or not responding with shock or no shock. The "omission" probability (p_{00}) is the probability of no response and no shock. The "punishment" probability (p_{11}) is the probability of both a response and a shock. The traditional avoidance contingency never omits shock on nonresponse trials ($p_{00} = 0$) and never presents shock on response trials ($p_{11} = 0$). Rats were trained on a discrete-trial paradigm with no intertrial interval. The first lever response changed an auditory stimulus for the remainder of the trial. Shocks were delivered only at the end of each trial cycle. After initial training under the traditional avoidance contingency, one group of rats experienced changes in omission probability ($p_{00} > 0$), holding punishment probability at zero. The second group of rats were studied under different punishment probability values ($p_{11} > 0$), holding omission probability at zero. Data from subjects in the omission group looked similar, showing graded decrements in responding with increasing probability of omission. These subjects approximately "matched" their nonresponse frequencies to the programmed probability of shock omission on nonresponse trials, producing a very low and approximately constant conditional probability of shock given no response. Subjects in the punishment group showed different sensitivity to increasing absolute punishment probability. Some subjects decreased responding to low values as punishment probability increased, while others continued to respond at substantial levels even when shock was inevitable on all trials (noncontingent shock schedule). These results confirm an asymmetry between two dimensions of partial avoidance contingencies. When the consequences of not responding included occasional omission of shock, all subjects showed graded sensitivity to changes in omission frequency. When the consequences of responding included occasional shock delivery, some subjects showed graded sensitivity to punishment frequency while others showed control by overall shock frequency as well.

One method of specifying the degree of dependence between responding and shock omission in avoidance procedures is through the conditional probability of shock delivery given a response or no response (Catania, 1971; Church, 1969; Gibbon, Berryman, & Thompson, 1974; Neffinger & Gibbon, 1975; Seligman, Maier, & Solomon, 1971). The traditional avoidance procedure in which responding (R) assures that shock (S^*) will not be delivered and not responding ($\sim R$) always results in shock delivery, has $P(S^*/R) = 0$ and $P(S^*/\sim R) = 1$. Partial avoidance contingencies result from varying these conditional probabilities away from 0 and 1. When the conditional probability of shock given a response is greater than 0, the result is partial punishment of avoidance responding. Conversely, when the conditional probability of shock given no response is less than 1, the result is partial reinforcement of not responding. When these two conditional probabilities are equal, a noncontingent (response-independent) shock schedule is in effect. For example, when $P(S^*/R) = P(S^*/\sim R) = .5$, subjects will receive 50% shock if they never respond and 50% shock if they always respond. Intermediate response levels will not alter this shock rate. The noncontingent procedure corresponds to zero correlation between response alternatives and reinforcement (Gibbon et al., 1974). When both conditional probabilities

---

[1]The work was supported by NSF Grants GB 34095 and BNS76-01229. The data reported were originally presented in the dissertation of Barbara L. Flye, in partial fulfillment of the requirements for the PhD degree, Columbia University, 1974. Reprints may be obtained from John Gibbon, New York State Psychiatric Institute, 722 West 168th Street, New York, New York 10032.

[2]Presently at New York Hospital/Cornell Medical College, Department of Psychiatry, 21 Bloomingdale Road, White Plains, New York.

are equal and high, shock density is high, while, when both are equal to zero, the traditional extinction procedure is in force.

Neffinger and Gibbon (1975), using a modification of a discrete trial procedure devised by Hineline and Herrnstein (1970), examined contingency and shock density effects by manipulating these two conditional probabilities. They found that some subjects were sensitive only to contingency without being sensitive to changes in shock density, whereas other subjects showed effects due to both of these variables. The contingency-sensitive-only subjects showed graded response levels with changes in contingency produced by either the punishment, $P(S^*/R)$, or partial reinforcement, $P(S^*/{\sim}R)$, variable, with responding decreasing to low levels as the noncontingent condition was approached. Other subjects showed this same pattern when studied with changing values of $P(S^*/{\sim}R)$, when $P(S^*/R)$ was kept at zero. However, some responding was maintained under the noncontingent schedules as long as shock density was not 0. For these subjects, responding appeared to reflect two sources of support: one an elicitation-like process which depended strictly on the frequency of shock delivery, and the other a contingency-sensitive process similar to that which controlled graded responding for the other subjects.

Avoidance contingencies specified by conditional probabilities are "ratio-like" in that the frequency of shock omission depends directly on the frequency of responding. The present experiments examined an "interval-like" method of specifying partial avoidance contingencies. In interval schedules, the rate of reinforcement is held approximately constant as long as responding is maintained above some minimum level. In a discrete trial context, rate of an event per unit time translates into the absolute probability of that event per trial. The partial avoidance contingencies studied here held the absolute probability of shock omission (and delivery) approximately constant.

Discrete trial avoidance contingencies are all referable to the basic 2 × 2 cross classification at the top of the right-hand column.

Columns represent omission or delivery of shocks, and rows represent the nonoccurrence or occurrence of responses. Entries in any cell represent the absolute probability of the conjunction of that row and that column. They are to be thought of as approximated over a long training period by the frequency of the conjoint events divided by the total number of trials. The rows accordingly sum to the probability of no response or response, respectively, and the column sums are the corresponding overall probability of no shock or shock. The conditional probability of shock given a response or nonresponse trial is shown below the tabulated material.

	$\sim S^*$	S^*	Σ
R	p_{10}	p_{11}	$P(R)$
$\sim R$	p_{00}	p_{01}	$P(\sim R)$
Σ	$P(\sim S^*)$	$P(S^*)$	1.0

$$P(S^*/R) = \frac{p_{11}}{p_{10} + p_{11}}$$

$$P(S^*/{\sim}R) = \frac{p_{01}}{p_{00} + p_{01}}$$

The space of all possible procedures represented by the four absolute probabilities is a three-dimensional tretrahedral volume, since the four absolute probabilities are subject to the single constraint that they must sum to 1.0 (Gibbon et al., 1974). Of the resulting three degrees of freedom, one is the response, nonresponse dimension controlled by the subject. The other two degrees of freedom specify the experimental contingency. The conditional probability specification represents one way of controlling these two degrees of freedom. The traditional avoidance contingency may be specified either by the conditional probabilities, $P(S^*/R) = 0$ and $P(S^*/{\sim}R) = 1$, or by the requirement that the absolute probability of no response and no shock, and the absolute probability of response and shock both equal zero, $p_{00} = p_{11} = 0$. When these two absolute probabilities are allowed to vary above zero, an alternative partial avoidance specification results.

The experiment reported here manipulated the omission probability, p_{00}, and the punishment probability, p_{11}, in two separate groups. In the Omission Group, the probability of omitting shock on nonresponse trials (p_{00}) was varied while the probability of shock on response trials was held at zero. For the Punishment Group, p_{11} was varied, with $p_{00} = 0$. In Figure 1, the space of possible procedures for the Omission Group is shown. Since punishment probability is fixed at 0, the resulting three absolute probabilities must sum to 1, and this restricts the space to one face of the tetrahedral volume. A similar triangular space has been studied recently by Rachlin and Burkhard (1978). The equilateral triangle is drawn so that the vertical dimension represents response probability. Response probability goes from 0 to 1, moving from the base to the top vertex. The right-to-left dimension shown below the triangle represents omission probability, p_{00}. The third probability, p_{01}, which must equal the remaining probability density, is shown above and to the right of the triangle. If omission probability is set at zero, then the traditional avoidance contingency results, in which not responding is always shocked and responding never shocked. In the triangular space, this is represented by the right edge, labeled A. Any response level is possible under the traditional avoidance procedure, and that corresponds to the fact that all horizontal lines intersect the $p_{00} = 0$ edge on the triangle. For omission probabilities greater than zero, all response levels are still possible. However, if response levels are greater than $1 - p_{00}$, subjects will not obtain all programmed omission events. This is precisely analogous to appetitive VI scheduling in which a minimum response rate is required to permit delivery of the programmed reinforcement rate. For example, on a VI 30-min schedule, a subject with very long interresponse times may experience a VI 60-min de-

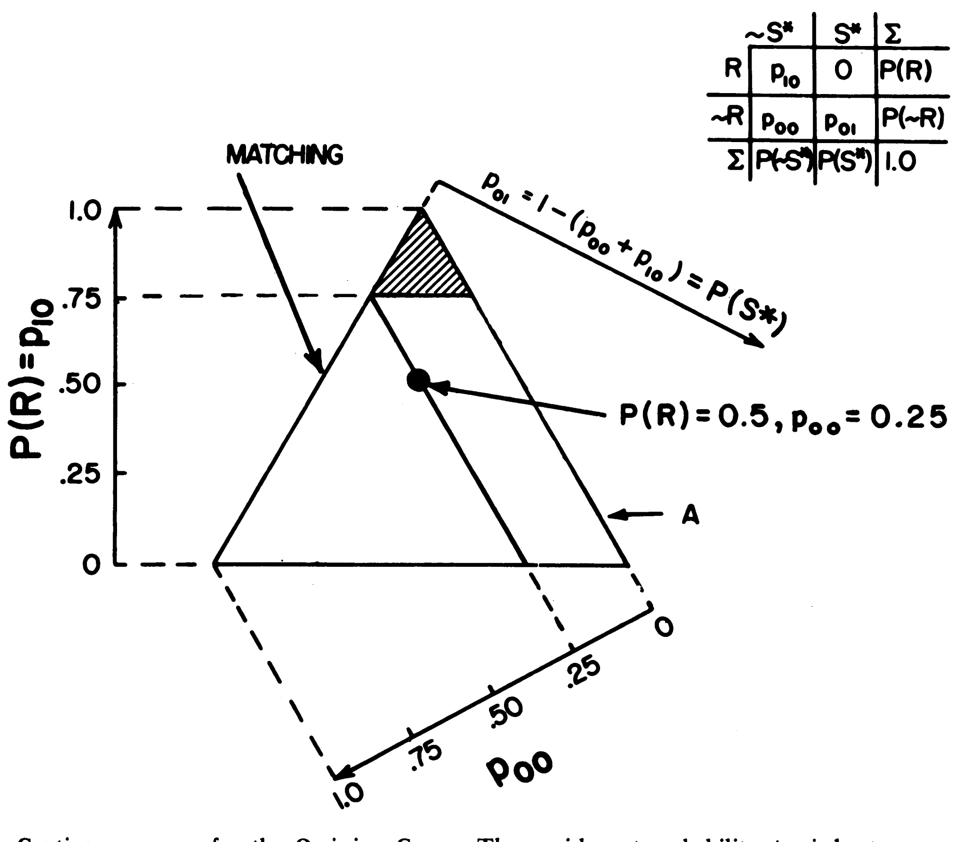

Fig. 1. Contingency space for the Omission Group. The punishment probability, p_{11}, is kept at zero and the remaining three probabilities must sum to 1. The corresponding dimensions are indicated by the arrows in the figure. They are: $P(R) = p_{10}$ (vertical dimension); the probability of a nonresponse and a shock, p_{01} (left to right dimension); and the omission probability, p_{00} (right to left dimension).

livery schedule. Thus, the absolute probability programming schedules are analogous to the appetitive VI schedules with the reinforcer being the omission of shock when no response occurs.

The rising line within the triangle and parallel to the right edge represents an omission probability of $p_{00} = .25$. Clearly, response levels above .75 (hatched area) preclude attainment of an absolute omission probability of this size. For response levels below .75, omission probability remains constant at .25 (the programmed value). A given response level results in a point along the rising line such as shown for $P(R) = .5$.

If responding closely approximated $1 - p_{00}$ —that is, if not responding matched the programmed omission value—the resulting points would lie along the left edge of the triangle labeled Matching. Such a strategy would result in the least shock (near zero) for the least effort (the minimum number of responses required to just exceed $1 - p_{00}$). Formally, this strategy might be viewed as a "relative reinforcement matching" function if reinforcement in the avoidance situation were construed to be the omission of shock on any trial. Then, reinforcement on response trials, represented by p_{10}, and reinforcement on nonresponse trials, represented by the programmed p_{00}, combine such that

$$P(R) = \frac{\text{Reinf. for } R}{\text{Reinf. for } R + \text{Reinf. for } {\sim}R} = \frac{p_{10}}{p_{10} + p_{00}} \quad (1)$$

because, if $P(R) = 1 - p_{00}$, $p_{10} + p_{00} = 1$. The matching edge, then, represents the minimum shock rate for the minimum response rate.

Noncontingent procedures in which shock delivery is independent of responding are only possible in this scheme when either all shocks are omitted, $p_{11} = p_{01} = 0$ (extinction), or when shock is delivered on every trial, $p_{10} = p_{00} = 0$.

The formal symmetry of the contingency table means that the space for variation in punishment probability looks identical to that in Figure 1, with p_{00} replaced by p_{11}, and p_{10} and p_{01} reversed.

The punishment and omission procedures differ with respect to a matching strategy as well. Under the omission procedures, subjects can "afford" to lower response rates and still receive no shock. Under the punishment procedures, responding is always the "best" strategy, since it is always the only way in which shock omission may occur. That is, if reinforcement is construed as shock omission, the matching equation above (1) has response probability = 1.0, independently of p_{11}, since p_{00} for the Punishment Group is maintained at 0. This is simply another way of saying that the "best" strategy under the punishment procedures is to respond on every trial since it is the *only* way in which shock omission may occur. Thus, while the contingencies are formally symmetric in the two groups, they are asymmetric from a matching of relative reinforcement point of view.

METHOD

Subjects

Eight naive male hooded (Long-Evans) rats and four naive male albino (Wistar) rats, housed in pairs and maintained on ad lib food and water, served. The subjects were selected as described below. They were approximately 60 days old at the start of experimentation.

Apparatus

Sessions were conducted in two standard operant conditioning chambers (Grason-Stadler Model #E3125D) housed in sound-attenuating enclosures. The chambers were modified by two stainless steel partitions that reduced the inside dimensions to a 13.33- by 13.33-cm square. This utilized 10 of the 18 stainless steel bar grids spaced 1.2 cm apart. The grids were connected to a shock generator (Grason-Stadler Model #E1064GS). Scrambled AC shock was delivered to each bar grid, the walls, and the response lever. The intensity used throughout the experiment was .8 ma; the duration was .5 sec.

A square retractable lever (Campden Instruments, Ltd., London, Model #C1-446) extended into the chamber 1.59 cm. It was mounted in the front wall of the chamber 8.89 cm above the floor. The lever required approximately 10 N of force to activate a Reed magnetic switch. The extension or retraction time was approximately .4 sec.

Auditory stimuli used in the experiment were a fixed frequency tone (1000 ± 10 Hz) and an adjustable clicker. The sound intensities in each chamber were matched using a sound level meter (General Radio Co., Type #1551A) so that each read 55 dB on the C scale.

Standard relay circuitry used to program the experimental conditions was located in an adjoining room. The occurrence of shock was programmed independently and individually for each experimental chamber. Randomized probability sequences of trial outcomes were punched into 16-mm film loops 100 units long. The film loop advanced with each trial, and an event which was not collected immediately was stored for delivery in a memory bank with a 25-unit capacity. Events were delivered when the appropriate row of the contingency table ($\sim R$ or R) was entered by the subject.

PROCEDURE

Basic Procedure

A discrete trial procedure (Hineline & Herrnstein, 1970; Neffinger & Gibbon, 1975) was used. The procedure when $p_{00} = 0$ and $p_{11} = 0$, the traditional avoidance contingency, is diagrammed in Figure 2. Sessions consisted of a series of trial cycles (top row) with no intertrial interval. Each trial was made up of a 20-sec response segment followed by a .5 reinforcement segment. Responses were effective only during the response segment, and when shocks were programmed they occurred only during the reinforcement segment. Each trial began with the onset of an auditory stimulus, either a 1000 Hz tone as shown here, or an 8-per-sec clicker, associated with no responding. The stimuli were counterbalanced across groups. If no response occurred within the response segment, the auditory stimulus was turned off and the lever retracted at the onset of the .5-sec reinforcement time. The first trial cycle in Figure 2 shows this sequence. At the beginning of the next trial, the lever was reextended into the chamber and the auditory signal turned on. Responses were not effective during the extension or retraction period. When a response occurred, the lever retracted and the auditory stimulus was changed from the tone to the click or vice versa. This is shown in the second trial cycle. At the end of the response segment, the post-response auditory stimulus was also turned off for the duration of the reinforcement segment. At the end of the reinforcement segment, the next trial began with the nonresponse signal and reextension of the lever into the chamber. Thus, each trial cycle began with the stimulus associated with no responding, which subjects could change to the other stimulus by making a response, and either stimulus was turned off during the final reinforcement period of each trial.

When the omission probability is greater than zero, a partial omission schedule is in effect. Figure 3 shows the scheduling procedure for three hypothetical examples: when responding does not occur for six trials (top diagram), when responding occurs on each of six trials (middle diagram), and when responding occurs on 50% of the trials (bottom diagram). The 50% omission schedule for each of these examples programs shock omission for a nonresponse occasion on Trials 3, 4, and 6. The top line of each diagram shows how omission events are stored in the memory bank for delivery on the next subsequent nonresponse trial. In the diagram on the top, no responses occur, and so the omission memory on Trials 3, 4, and 6 reverts to zero with the nonoccurrence of shock. In the middle diagram, responses, indicated by circled x's on the response lever line, occur on every trial. The omission events, therefore, may not

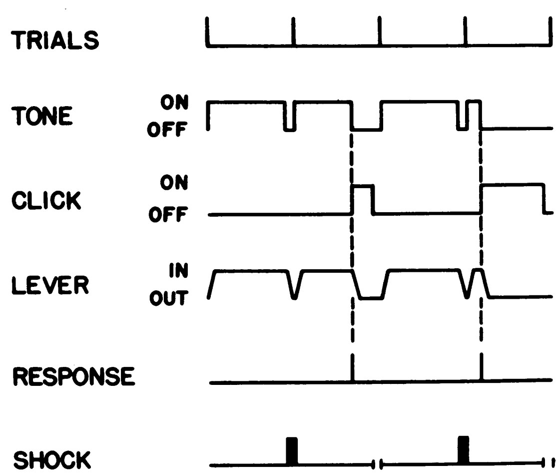

Fig. 2. Experimental paradigm. Trials are indicated on the top row and follow each other with no intertrial interval. They are composed of a response segment during which a tone or a clicker stimulus is on, followed by a brief reinforcement segment at the end of the trial cycle. When shock occurs (bottom row), it occurs only in the reinforcement segment. Trial cycles begin with lever extension and the onset of a tone associated with not responding. Responses (second cycle) change the auditory stimulus to a click and retract the lever for the remainder of the trial cycle.

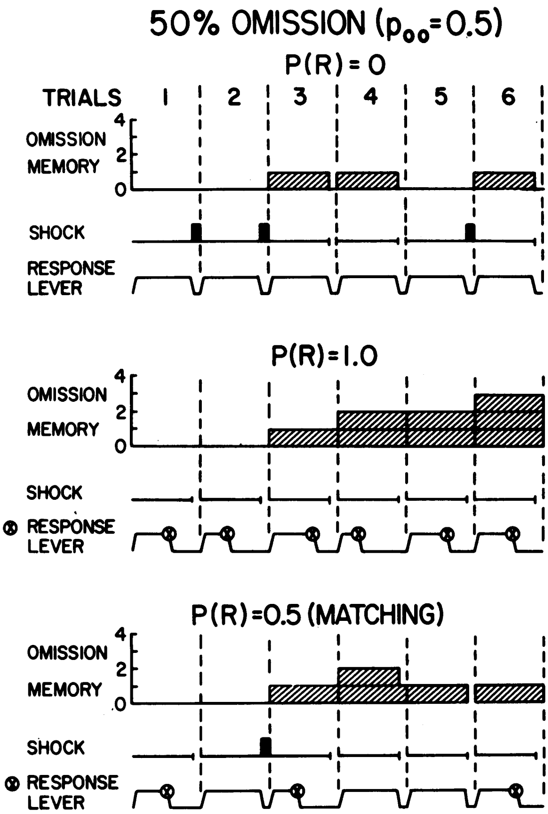

Fig. 3. The absolute programming procedure with the "Omission Memory" for $p_{oo} = .50$. Three different response frequencies and their outcomes are indicated in the top [$P(R) = 0$], middle [$P(R) = 1.0$], and bottom [$P(R) = .5$] diagrams. See text for additional details.

be delivered on these six trials and are stored successively in the memory. In this example, should the subject fail to respond on three successive trials after the six response trials, no shock would be delivered.

The bottom diagram shows what would occur if responses were made on Trials 1, 3, and 6. A shock is received on Trial 2 because no omission event is programmed. On Trial 3, nonresponse omission is programmed but a response is made, thus preventing its delivery. On Trial 4, a second omission event is stored in memory and one of these is delivered at the end of Trial 4, since no response occurs. Failure to respond in Trial 5 results in the delivery of the second omission event, so that the memory system reverts to zero. On Trial 6, another omission event is programmed, but the response on Trial 6 forestalls its delivery.

In these three examples, omission probability is held constant, but the conditional probability of shock given no response depends on the subject. When no responses are made, the conditional probability of shock given no response is .5, the same as the omission probability. In the middle panel, responding occurs on all trials, and so conditional probability of shock given no response is not defined. In the bottom panel, the conditional probability of shock given no response is 1/3. Thus, changes in response levels alter the conditional probability of shock given no response; but over a long series of trials, the 50% omission schedule results in shock omission on 50% of the nonresponse occasions, provided at least 50% of the trials are nonresponse trials. On the omission schedule, when subjects respond on every trial, they guarantee no shock, but for maximum response effort. When subjects respond on slightly more than 50% of the trials, they receive very few shocks for minimum response effort. When subjects respond on less than 50% of the trials, the overall shock rate is intermediate between .0 and .5, but the frequency of shock on nonresponse trials remains constant at .5.

The study was conducted in two replications. The first two groups of rats (hooded) were studied using the paradigm in Figure 2. For the second two groups (albinos), the lever was continuously available (extended). This change was made to determine if lever movement contributed to responding, particularly under punishment. No differences were observed, however, between the first and second replications.

Latencies to respond in each trial were recorded in twelfths of the interval. Sessions were conducted 6 days a week and lasted for 200 trials, with the first 50 trials ("warm-up") deleted from the data analysis. These data were collected and analyzed separately in anticipation of possible effects. Warm-up effects were not present.

Subject Selection

Subjects were selected for the experiment using a shaping procedure followed by training under the traditional avoidance contingency. During the shaping sessions, shocks were delivered unsystematically at a rate of approximately one every 5 sec. A response by the animal meeting the shaper's criterion

Table 1

Order of schedule presentation. H- and A- prefixes indicated hooded and albino subjects. The footnotes below the table give the number of days at a schedule value when exposure exceeded 9 days. EXT means extinction ($p_{11} = p_{01} = 0$), and *N-C S** means noncontingent shock ($p_{00} = p_{10} = 0$).

Omission Group							
		p_{00} values					
		H-35	*H-136*	*H-145*	*H-156*	*A-14*	*A-4*
Condition	1	.50	.50	.50	.50	.50	.50
order	2	EXT	EXT	.85[d]	.85[b]	.50	.50
	3	.85	.85	EXT[d]	EXT[c]	.85	.85
	4	.25	.25			.95	EXT
	5	.50	.50			.25	.25
	6					EXT	.95
	7					.50	.50
	p_{11} =	*N-C S**	*N-C S**	*N-C S**	*N-C S**	*N-C S**	*N-C S**
						*N-C S**	*N-CS**[b]
Punishment Group							
		p_{11} values					
		H-46	*H-56*	*H-126*	*H-134*	*A-156*	*A-234*
Condition	1	.05	.05	.05	.05	.05	.05
order	2	.25	.25	.25	.25	.05	.05
	3	*N-C S**	*N-C S**	*N-C S**	*N-C S**	.25	.25
	4	*N-CS**[c]	.85	*N-C S**	*N-C S**	.25	.25
	5	.25	.50		.25	.10	.10
	6		.25		.50[a]	.20	.20
	7		.50		*N-C S**	.50	.50
	8					*N-C S**	*N-C S**
	9					*N-C S**	*N-C S**

[a]12 days.
[b]15 days.
[c]18 days.
[d]21 days.

had the following consequences: The auditory stimulus changed, the train of shocks was interrupted for 26 sec, and (for the hooded subjects) the lever retracted from the chamber. Twelve hooded and 29 albino rats experienced 2 shaping sessions and were then continued for 9 days of traditional avoidance contingency training, on the basic paradigm (Figure 2). Eight hooded and four albino subjects acquired response probabilities above .60 after 9 days, and these animals served in the experiment proper.

Experimental Procedure

After selection, all subjects were continued on the traditional avoidance contingency for 18 days. After this pretraining, subjects were divided into the Omission Group and Punishment Group, with four hooded and two albino animals in each, matched approximately for response probability on the last 6 days of pretraining. They were then studied at a sequence of p_{00} or p_{11} values with a redetermination at the traditional avoidance contingency condition between each point. Determinations at omission and punishment values and recovery under the traditional avoidance contingency were continued for a minimum of 9 days and extended beyond that time if the final 3 days varied by more than 10% from the 3-day mean. The sequence of schedule values for both groups is shown in Table 1. Four subjects in the Omission Group were studied at $p_{00} = .5$ after training on other values to observe possible order effects. After training at the p_{00} values shown, all subjects in the Omission Group were studied under the 100% noncontingent shock condition (N-C S*; $p_{00} = p_{10} = 0$).

In the Punishment Group, the p_{11} values used depended on the performance of individual subjects. Two subjects were relatively insensitive to the punishment variable and were not studied at as many values as the others. Variations in response patterns and trial outcome are symmetrical to those diagrammed in Figure 3.

RESULTS

Omission Group

Response probability pooled over three session blocks is shown for three omission subjects in Figure 4, with the p_{00} values indicated above each panel. Recovery determinations at the traditional avoidance contingency (labeled A) intervene between each determination. The data for H-136 shows orderly decrements in responding with increasing p_{00} values. Recovery at the traditional avoidance contingency condition was relatively rapid and to high response probability levels. The decrements in response probability also occurred quite rapidly, with response adjustment virtually complete within the first three days. Note that the $p_{00} = .5$ redetermination produced response levels quite similar to the earlier treatment. Data from the three subjects not shown in Figure 4 were comparable to H-136, showing orderly and rapid decrements in responding with increases in p_{00}.

Subject H-145 showed somewhat more protracted adjustment of response levels to different p_{00} values, though, again, the response

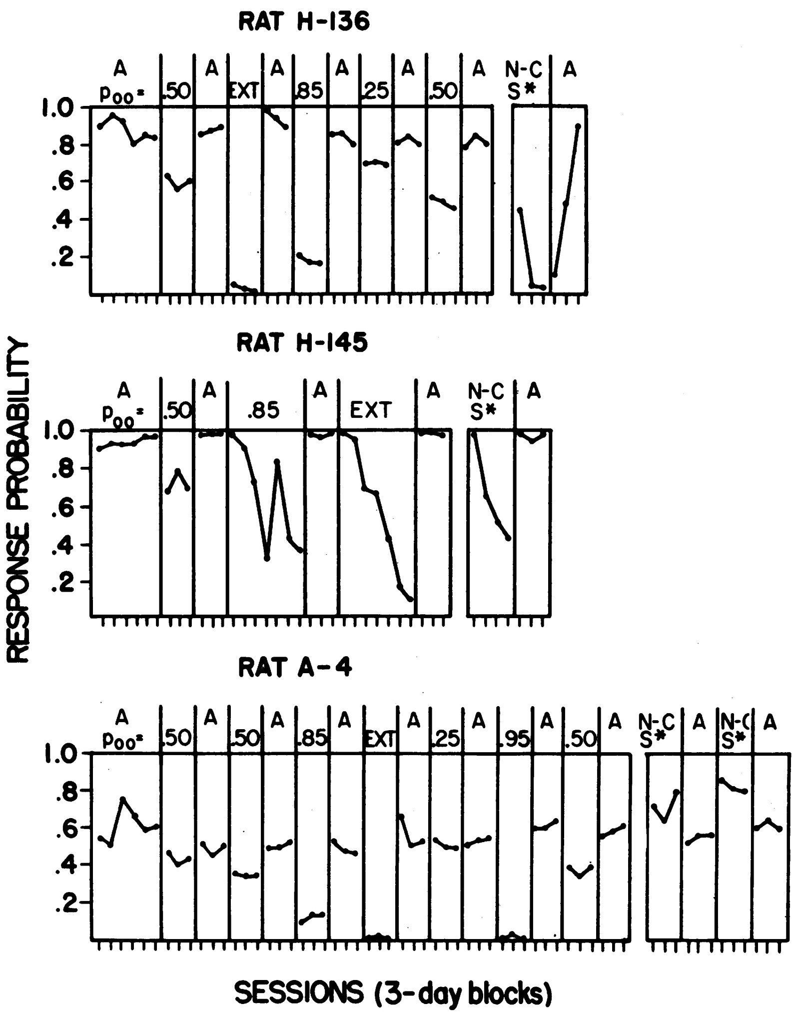

Fig. 4. Response probability pooled over three session blocks under successive p_{00} values for three subjects in the Omission Group. Redeterminations at the traditional avoidance contingency intervene between each treatment value. The panel to the right of the break shows responding under noncontingent shock.

levels attained are monotone with the omission probability.

Subject A-4 had the poorest avoidance performance under the traditional avoidance contingency of any subject in the group (about .55). However, orderly decrements with increasing p_{00} were also observed for this subject, and the redetermination of $p_{00} = .5$ again replicated the earlier levels.

The panels to the right of the break for each subject show responding under the 100% noncontingent shock schedule (*N-C S**). H-136 stopped responding entirely under noncontingent shock, just as it did under the traditional extinction condition. H-145 showed substantial responding under noncontingent shock, with behavior stabilizing at about $P(R) = .4$, in contrast to the near-zero responding of this subject under extinction. A-4 showed the highest response level of any subject under noncontingent shock. For this subject, shock on every trial supported more responding (about .8) than the shock omission contingency of the traditional avoidance contingency (about .55). The other subjects showed some variation between these extremes of substantial responding and no responding under the noncontingent 100% shock schedule.

The functional relation between response probability and programmed omission values is shown in Figure 5 for each subject. The data are taken from the last three days at each determination. Values shown for $p_{00} = .5$ are averages of original and redetermined points. All subjects' data are somewhat below 100% responding under the baseline avoidance condition ($p_{00} = 0$) and slightly above no responding at the traditional extinction point. Their responding between these values was a smooth monotone decreasing function of p_{00}. The matching relation, in which responding produces approximately zero shock for the least effort, is the dashed line with slope -1. Response levels lie slightly above this line for most of the omission values studied. This approximate matching means that subjects were

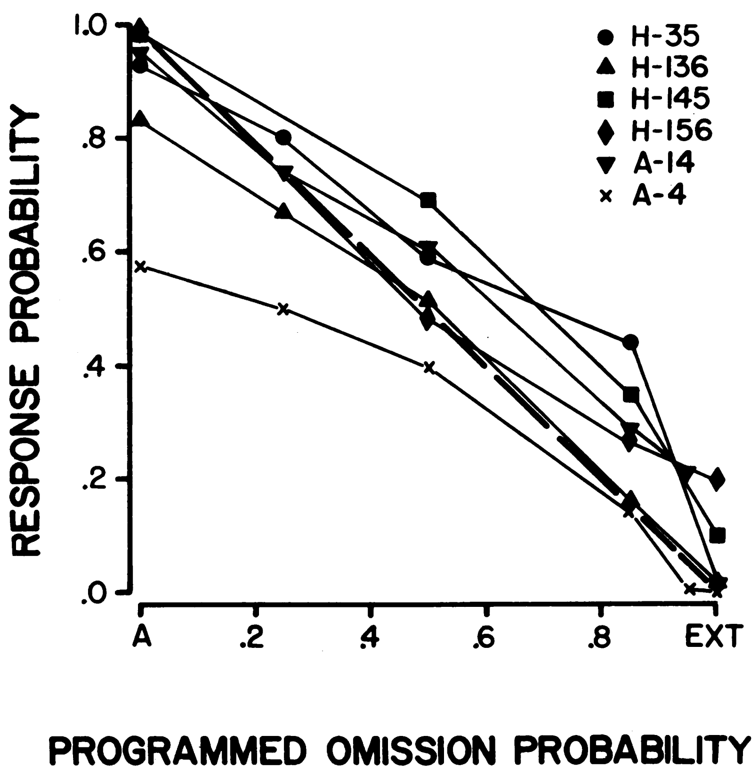

Fig. 5. Response probability as a function of the programmed omission probability. Points at the extreme left represent a responding under the traditional avoidance contingency and those at the extreme right under extinction. The heavy dashed line, with slope -1, represents the matching relation in which nonresponse levels just match the programmed omission probability.

responsive to the "unlimited hold" feature of the memory bank for programming omission. They "used up" the stored shock omission events (cf. Figure 3) at a rate just slightly lower than a nonresponse frequency which would result in substantial shock. Rat A-4, in contrast, responded consistently below the matching level throughout the range of programmed p_{00} values. While response probability decreased with increasing omission probability for this subject also, it did so too rapidly to result in the low shock rates characteristic of the other subjects throughout most of the p_{00} range.

The programmed omission probability and the behavior of the subjects interacted in the above manner to produce a low and relatively constant conditional probability of shock given no response; Table 2 presents this conditional probability for the points shown in Figure 5. Fourteen of the 18 values are below 5%. Most of the subjects achieved this low shock rate on nonresponse trials for all p_{00} values greater than or equal to .50. The exception was again with Rat A-4, which achieved the 5% rate only at the highest p_{00} values. It is noteworthy that 5 to 10% is just the level found by Boren and Sidman (1957) and Neffinger and Gibbon (1975) at which responding was just maintained when the $P(S^*/{\sim}R)$ was programmed experimentally.

Latency data for all subjects was computed as the conditional probability of a response latency of t sec given a latency of at least t sec (latency per opportunity). Functions from selected conditions for each subject are shown in Figure 6. The data are based on response trials only, and points were not computed unless the number of opportunities exceeded 20. The left-most panels show response probability functions under the avoidance baseline condition. Two functions are shown, one for all latencies (circles) and another for latencies after a shocked (nonresponse) trial (triangles). For three subjects response probability was so high that complete postshock functions were not available. H-35 and H-145 showed substantially flat conditional probability of response over the duration of the trial. A flat conditional probability function (corresponding to an exponential relative frequency function) indicates no temporal control. The other subjects showed some rise in response probability as the end of the preshock interval was approached. The postshock functions, if anything, show a somewhat sharper temporal discrimination than the overall latencies, with response strength increasing as the preshock interval elapses.

The next column shows latency distributions collected over the last three days at the 50% omission schedule. These functions are similar to those obtained under the baseline condition, and the second function, for responding after omission trials (x's), is not different from responding overall. It is noteworthy that the two subjects showing little temporal control under the avoidance baseline, H-35 and H-145, exhibited some temporal control under the $p_{00} = .5$ schedule.

Three subjects exhibited sufficient responding under the noncontingent shock condition to allow a latency distribution analysis. Where possible, latencies after a nonresponse shock (triangles) and after a response-produced (punishment) shock (squares) are presented separately. It is clear that the function forms are very similar to each other and to the overall function. For the subjects for which response probability was sufficiently high to observe responding throughout most of the latency range, temporal control did not appear to be substantially different late in the preshock interval under noncontingent shock than it was under the previous omission schedule. A-4 again represents an exception to this characterization inasmuch as it showed little evidence of temporal control under noncontingent shock. The other albino subject, A-14, showed a high response probability in the first category under noncontingent shock, similar to its postshock function under the traditional avoidance contingency. For the albino subjects, the lever was continuously

Table 2

The conditional probability of shock given no response at each programmed p_{00} value for each Omission Group subject. Data are taken from the last three days on each condition.

	$p_{00} =$			
Subject	*.25*	*.50*	*.85*	*.95*
H-35	.038	.021	.008	
H-136	.202	.045	.011	
H-145		.004	.000	
H-156		.026	.015	
A-14	.110	.039	.006	.006
A-4	.471	.181	.040	.047

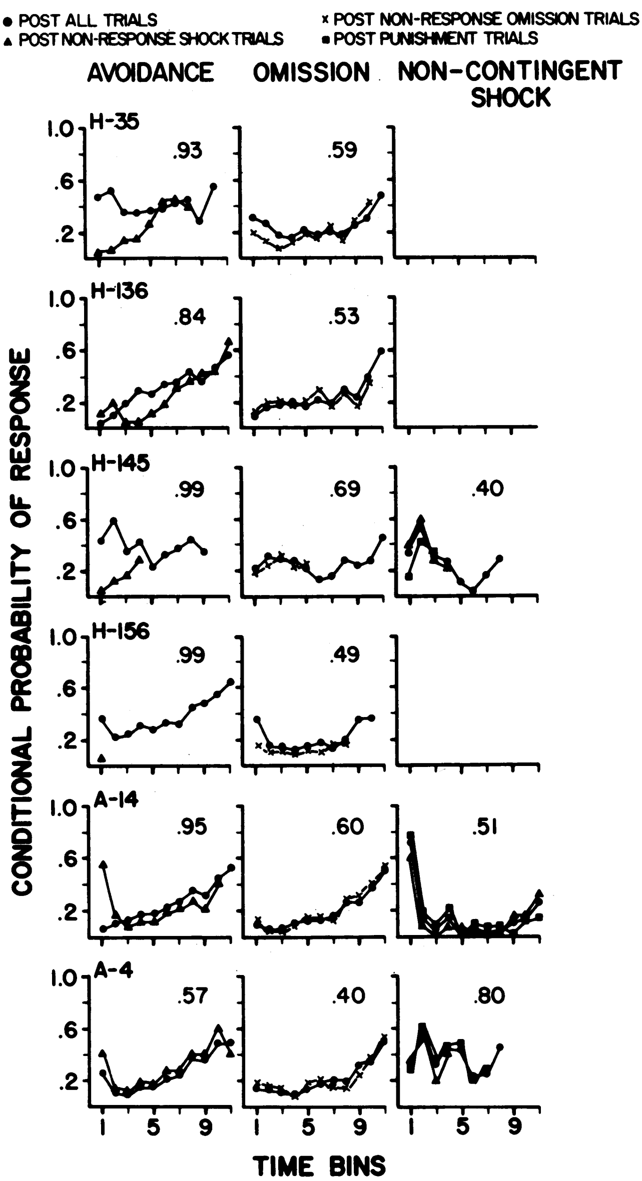

Fig. 6. Conditional probability of a latency of t sec given a latency of at least that long. Each row represents a subject. Performance under the traditional avoidance contingency is shown in the left-most panels, under the 50% omission schedule in the middle panels, and under noncontingent shock in the right-most panels. The inset value is response probability. Different functions correspond to responding following different preceding trial events, indicated in the legend.

available, so that this pattern may reflect postshock "burst" (Boren, 1961). It is noteworthy, however, that A-14 responded less under noncontingent shock than A-4, which did not exhibit this pattern.

Sequential Analysis

A pattern of responding which might produce something like the smooth reduction in response levels with increasing omission probability observed here is one in which subjects respond for a time directly after shock, and then cease responding until another shock is delivered. Such a pattern would produce a reduction in responding with increasing omission probability since longer waits for shock would result from larger programmed omission probabilities. Data from the 50% omission schedule of two hooded and two albino subjects were analyzed for sequential patterns of this sort. The data were taken from all determinations at the 50% omission schedule value, which consisted of 9 sessions for the albinos and 6 or 3 sessions for the hoodeds. Response and nonresponse occurrences were subjected to a run-length analysis, and frequency distributions of the number of responses in a row and the number of nonresponses in a row were calculated for each subject. Three subjects had very low shock rates on some of the sessions, and the two hoodeds received no shocks whatever on some sessions. The warm-up data (first 50 trials) were included in this analysis, which means a total of 250 trials per session. Run length distributions for sessions containing shocks (filled circles) and sessions containing no shocks or few shocks (open circles) are shown for these subjects in Figure 7. Numbers in parentheses on the right are the number of sessions represented in each function. The data on the left are for responses in a row and those on the right for nonresponses in a row. Points indicated by x's with numbers next to them are run lengths with frequencies less than 10. Frequencies below 10 were included only when the distribution would otherwise consist of one or two points. The ordinate is a log scale, and each subject's data has been displaced vertically by one log unit. The ordinate values represent the relative frequency (probability) of a given run length. On a semilog plot of this sort, a geometric distribution form becomes linear, and it is clear that that is a reasonable characterization of these data. The frequency of successive responses (or nonresponses) in a row decreased geometrically with run length in the manner characteristic of a random system with a constant probability of response (or nonresponse). Moreover, sessions with no shocks or few shocks (H-35, H-156, and A-14) showed precisely the same character, and overall response and nonresponse levels are comparable to shock sessions. Thus it seems that subjects adjusted response levels over a substantial period of time rather than on a trial-by-trial basis. Response probability remained approximately constant from trial to trial and session to session, and was about the same when there were no shocks in a session as when there were some shocks in a session. These data argue that, whatever the mechanism for adjusting response probability downward with increases in omission probability, it cannot depend critically on local changes when shocks occur. Nonresponse trials were not followed by a wait to the next shock; rather, the probability of terminating a nonresponse run remained constant whether shocks had been experienced or not.

Postshock responding was examined in another analysis cross-classifying response and nonresponse trials immediately following shocked and unshocked trials. Data from the four subjects in Figure 7 are shown in Table 3. Chi-square values for association between a shock on Trial N and response on Trial $N+1$ appear next to each fourfold table. Except for Rat A-4, none of the values are significant ($p > .10$), indicating that response probability on a given trial did not reflect the presence or absence of shock on the preceding trial. In contrast, the chi-square value for A-4 is highly significant, indicating a clear tendency toward more responding just after shock. Shock rates for this subject were high (30 to 40 per session) in the data of Figure 7, precluding a run-length analysis for shock-free data. Thus, A-4 was unusual in showing (a) the lowest avoidance baseline, (b) the highest noncontingent shock responding, and (c) a local dependence of responding on previous shock.

The sequential findings suggest path independence at asymptote under the omission schedules. They corroborate the lack of sequential effects in latency (Figure 6). Subjects evidently adopt a response probability appropriate to the current omission probability but are insensitive to local changes in the pre-

Table 3

Response and nonresponse frequencies following shocked and unshocked trials under the 50% omission schedule.

Subject	Trial *N*	Trial *N+1* ~*R*	*R*	χ^2	
H-35	*S**	3	7		
	~*S**	428	762	.004	*NS*
H-156	*S**	1	6		
	~*S**	244	349	1.10	*NS*
A-14	*S**	16	37		
	~*S**	692	1055	1.54	*NS*
A-4	*S**	142	154	49.3	*S*
	~*S**	1043	461		$P < 0.001$

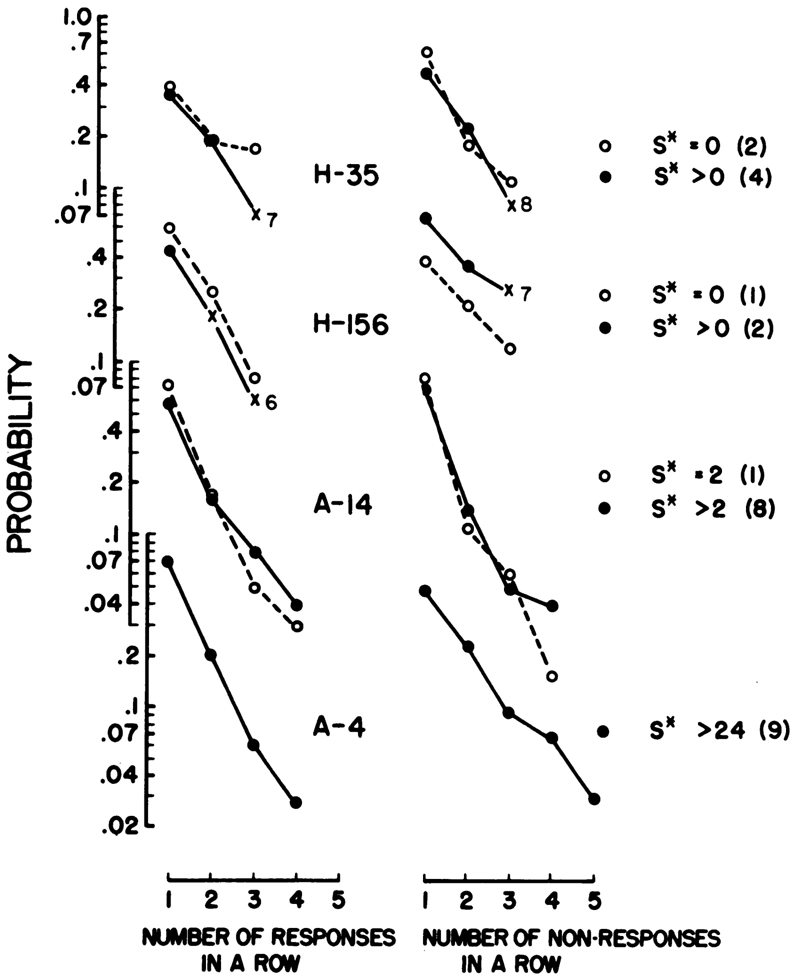

Fig. 7. Run-length distributions for responses in a row (left column) and nonresponses in a row (right column) for subjects under the 50% omission schedule. Sessions with no shocks or few shocks are indicated by open circles and shock sessions by closed circles. The ordinate is a log scale, and a straight line function indicates a geometric (constant probability) run-length distribution. The number of sessions included in each function is indicated in parentheses in the legend.

vious response, nonresponse, or shock occurrence.

Punishment Group

Two subjects, H-46 and H-126, were insensitive to the punishment variable and were studied at only a few p_{11} values. A third subject, A-234, was highly sensitive to p_{11}, showing nearly complete suppression for all p_{11} values greater than .25. The other three subjects showed strong order effects. Their data are shown in Figure 8. Response probability is shown in three-session blocks and alternate panels denote successive p_{11} treatments. Consider H-56 first. Increasing p_{11} values produced no change until the first treatment at the noncontingent condition, when shock occurred on every trial. This initial exposure resulted in about 30% responding. Subsequent exposure to p_{11} values of .85 and .50 produced similar low response levels. The next treatment at $p_{11} = .25$ led to a substantial *reduction* in response level from the previous .50 treatment, in contrast to the first exposure to $p_{11} = .25$ which resulted in no decrement whatever. When the punishment variable was again set at $p_{11} = .50$, responding was virtually eliminated. Extended exposure to punishment

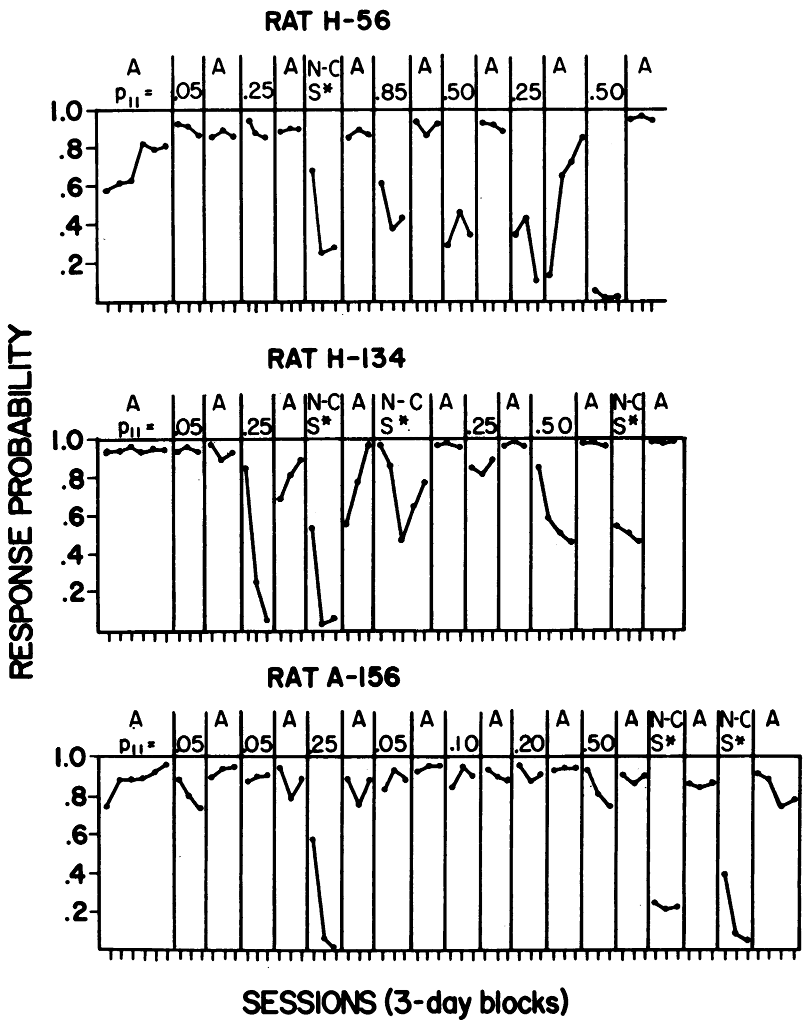

Fig. 8. Response probability pooled over three-session blocks under successive p_{11} values for three punishment subjects. Redeterminations at the traditional avoidance contingency intervene between each punishment value.

evidently "sensitized" this subject and resulted in increased suppression that was not reversible.

Rats H-134 and A-156 showed just the opposite effect of order of exposure. For both of these subjects, the first exposure to $p_{11} = .25$ resulted in dramatic reductions in responding down to near zero levels. However, for both subjects, later exposure to similar or higher punishment values (note performance at $p_{11} = .50$) resulted in less suppression than the initial exposure to $p_{11} = .25$. Thus, for these subjects, early exposure showed sensitivity to the punishment variable, whereas later exposure showed what might be called adaptation to punishment with less suppression produced at comparable and larger punishment values.

Figure 9 presents response probability as a function of the programmed punishment probability for all subjects in the Punishment Group. The traditional avoidance contingency is represented on the left ($p_{11} = 0$) and the noncontingent shock schedule on the right (*N-C S**). Three subjects, H-46, H-126, and A-234, have one function shown in the figure. These subjects did not change their sensitivities with increasing exposure to the punishment schedule. The three other subjects in the group have two functions plotted. The filled symbols connected by solid lines are for the first determination, and the open symbols connected by dashed lines are for a later determination. The values in the later determination were taken from exposures beyond the point at which a clear change in sensitivity had taken place (H-56 under $p_{11} = .5$, H-134 under $p_{11} = .10$, and A-156 under $p_{11} = .10$). Redetermination values within the early and late functions are represented by averages. It is clear from these data that subjects exhibited either increasing sensitivity with order of exposure (H-56) or decreasing sensitivity with increasing exposure (H-134 and A-156). However, the functions with a steep slope indicating sensitivity to punishment look quite similar to each other, as do those showing less sensitivity with a shallower drop in response levels.

The rising diagonal dashed line with a slope of 1.0 represents the analogue for the Punishment Group of the matching line for the Omission Group. Below this line increases in the programmed probability of punishment represent no change in the experienced shock rate, since shocks occur on virtually every trial for response levels in this range. For example, when punishment is programmed at $p_{11} = .25$, response levels below one quarter of the total number of trials require that virtually all trials must be shocked since the programmed p_{11} value is greater than the frequency or response-plus-shock trials which can be delivered. Further increases in p_{11} values do not result in graded response levels, since once responding falls below the diagonal, subjects experience shock on every trial. This is reflected in the flat character of the functions for points to the right of the diagonal.

Above the diagonal there is no clear representative form for these functions. Some subjects showed rapid suppression with increasing punishment, either early or late in training, while others showed more gradual decreases in responding as punishment levels increased. The unconnected points to the right of the figure represent noncontingent shock performances for Omission Group subjects studied under this condition after training

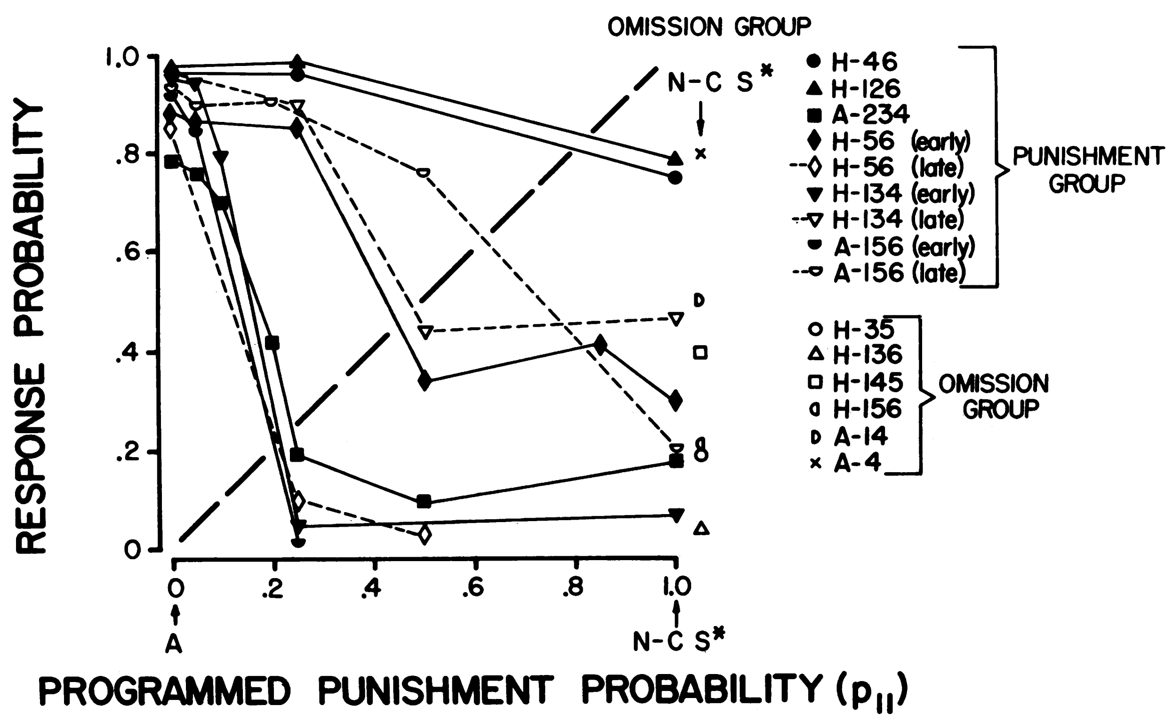

Fig. 9. Response probability as a function of the programmed punishment probability. The traditional avoidance contingency is on the extreme left and the noncontingent shock condition on the extreme right. Noncontingent responding of the omission group is shown to the right of the functions for the Punishment Group. For three subjects which showed changing sensitivity to the punishment variable with increasing exposure, two functions, early and late, are shown. Points below the heavy dashed diagonal line represent response probabilities too low to result in attainment of the programmed punishment probability. Thus, for all of these points, shocks occur on virtually every trial (noncontingent shock).

on the omission schedules. They demonstrated a broad range of suppression levels comparable to the range observed with the punishment subjects. Only 2 of the 12 subjects (H-136 and H-134) showed very low response levels under noncontingent shock. These subjects responded at high levels under the avoidance baseline condition, and thus may be thought of as "contingency sensitive" only (Neffinger & Gibbon, 1975). The other subjects show a range of response levels supported by shock alone. Thus these subjects' data reflect shared control between shock density and the shock omission contingency.

Latency distributions for the punishment animals are shown in Figure 10. The measure again is the conditional probability of a response given a latency of at least the abscissa value. The left column represents avoidance baseline training. The two subjects which were relatively insensitive to the punishment variable are shown in the top two rows. Except for H-126, all subjects showed some temporal discrimination of shock delivery time in their overall latencies (circles). Responding after a nonresponse plus a shock (triangles) showed a somewhat sharper temporal gradient with a lowered response probability in the early or middle portion of the trial.

The middle column presents latency functions obtained under $p_{11} = .25$ for the hooded subjects and $p_{11} = .20$ for the albinos. Data were taken from the determinations with the most responding. Several features of these distributions differ from those observed under the traditional avoidance contingency. Subjects tended to show a U-shaped function form for overall responding (circles). Early portions of the trial were associated with moderate response probability, followed by a later decline during the middle range of latencies, followed in turn by a rise later in the trial. The early high levels of response probability were not seen after shock trials (either response or nonresponse produced). Where possible, punishment trials and nonresponse trials are shown separately. These functions begin at low levels and show similar temporal control to the functions for responding after a nonresponse trial under baseline training.

The right-most panels represent latency distributions under the noncontingent shock schedule. Distributions are not shown for determinations with less than 25% response

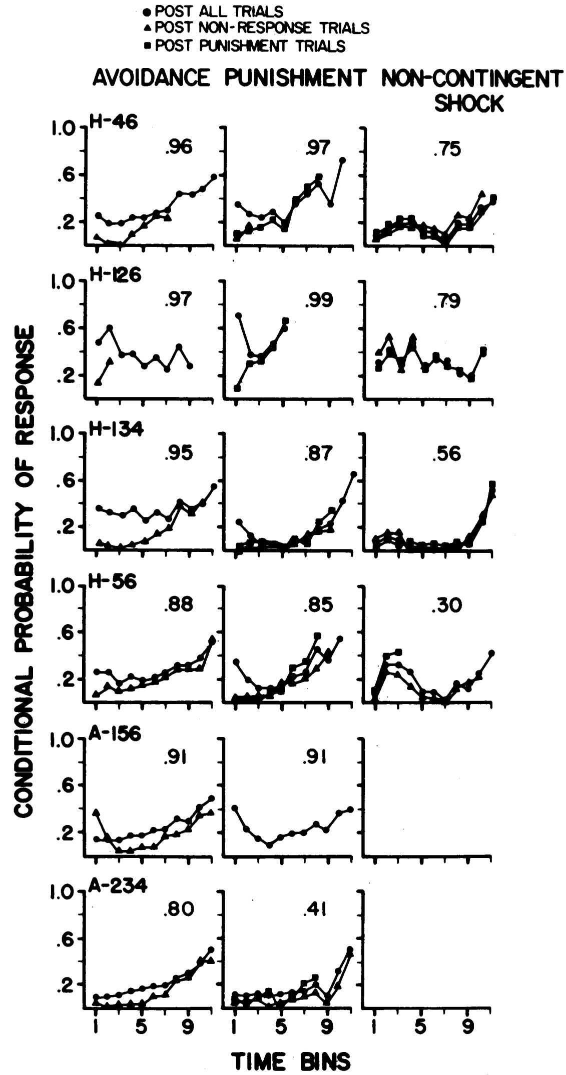

Fig. 10. Conditional probability of a latency of t sec given a latency of at least that long. Each row represents one subject from the Punishment Group. Performance under the traditional avoidance contingency is shown in the left-most panels, under a partial punishment condition in the middle panels, and under noncontingent shock in the right-most panel. The inset value is response probability, and the noncontingent shock distributions were not computed unless this value was greater than .25. Different functions within each panel correspond to response probability following different preceding trial events, indicated in the legend.

levels. Where possible the data were again separated into latencies following response trials and nonresponse trials, and for all subjects these look very much alike, and like the function for overall latency. With the exception of H-56, the function forms are similar to those observed under the traditional avoidance contingency and partial punishment. Three of the subjects showed temporal control at about the same level of precision observed earlier (H-46, H-134, and H-56), and H-126 showed the same lack of temporal control observed earlier. One subject (H-56) showed some tendency to respond at early and intermediate time periods in the trial at somewhat higher levels than observed earlier, but the later portion of its functions show the same degree of temporal control as this subject evinced in the previous treatments. A postshock "bursting" pattern was not clearly present for these subjects, even the albinos which had the lever available immediately after shock.

DISCUSSION

The present results parallel and extend the findings of partial avoidance contingencies defined on conditional probabilities reported by Neffinger and Gibbon (1975). The conditional probability of shock delivery, given a nonresponse instance, may be thought of as a shock omission variable in that lowering that probability results perforce in an above-zero frequency of nonresponse trials which end in no shock. Similarly, varying the conditional probability of shock given a response results in above-zero frequencies of response-plus-shock trials. Thus, the conditional probability schedules and the absolute probability schedules are close relatives, and the present results will be discussed in the light of the corresponding manipulation in conditional probability terms.

Omission

Neffinger and Gibbon found that all subjects eventually ceased responding as the traditional extinction condition $[P(S^*/R) = P(S^*/\sim R) = 0]$, was approached. However, the functions for $P(S^*/\sim R)$ were extremely steep ones with little reduction in responding observed until the conditional probability of shock given a nonresponse trial was reduced to about .05. This is the level reported by Boren and Sidman (1957) as that required to maintain behavior in a free-operant avoidance schedule. Those results contrast with the graded responding observed in the present experiment throughout the range of the shock-omission variable. All subjects showed a graded, nearly linear response to the increasing probability of omission. The difference is an important one and rests on a fundamental difference in the manner of experimental control of the partial schedules. When subjects are exposed to the absolute probability schedule with the memory feature (Figure 3), response frequencies modulate the conditional probability of shock given no response while absolute omission probability is held constant. In the conditional probability specification, the reverse is true. Subjects may produce high or low absolute probabilities of shock per session, but the conditional probability of shock given a nonresponse instance is maintained constant by the experimenter. The result is that when the conditional probability of shock given nonresponse is decreased experimentally, a constant level of responding results in fewer shocks or, equivalently, an increase in the absolute frequency of nonresponse plus no-shock. This reciprocal relationship may be displayed formally as

$$p_{00} = [1 - P(R)][1 - P(S^*/\sim R)], \qquad (2)$$

which follows from the contingency square specification (Gibbon et al., 1974). The rightmost term is equal to the conditional probability of no shock given no response. This variable is proportional to the absolute probability of omission, p_{00}, with the proportionality parameter being the probability of a nonresponse occurrence. This relation may be recast in a manner which allows comparison between the conditional and absolute probability manipulations. Equation 2 is precise only in case the absolute omission probability, and the conditional probability of shock given no response, are those actually experienced by the subject. Calling these obtained probabilities $\bar{p}_{00}$ and $\bar{P}(S^*/\sim R)$, response probability may be written

$$P(R) = 1 - [1 - \bar{P}(S^*/\sim R)]^{-1}\, \bar{p}_{00}. \qquad (3)$$

In this form the discrepancy between conditional probability and absolute probability findings may be resolved.

Imagine that $P(S^*/\sim R) = .05$ is an approximate "threshold" for avoidance responding. Under the absolute probability schedule, when the obtained $P(S^*/\sim R)$ rises above this value, responding occurs for, say, several trials. This drives $P(S^*/{\sim}R)$ below threshold, so that responding stops, which in turn increases $P(S^*/{\sim}R)$, and so on. The result is then a fluctuation of $P(S^*/{\sim}R)$ around .05. Equation 3 then has

$$P(R) \simeq 1 - (1.052)\bar{p}_{00}, \quad (4)$$

the approximate matching observed in Figure 5.

This point may be demonstrated by comparing performances from the two experiments on an equivalent measure, namely, the absolute omission probability actually obtained, $\bar{p}_{00}$. This is done in Figure 11. The Neffinger and Gibbon data are indicated by open symbols and the absolute probability data from the present experiment by filled symbols. The left-most point for $\bar{p}_{00} = 0$ represents the traditional avoidance contingency in which all nonresponse occurrences are shocked. Under this avoidance baseline condition, subjects did not all perform equally efficiently. However, as omission probability increases, the data tend to approximate more closely the matching line $1 - \bar{p}_{00}$. Data from the Neffinger and Gibbon study is largely restricted to low values of $\bar{p}_{00}$, since responding remains maximal for these subjects throughout most of the range of conditional probability studied. Even over this restricted range, however, it is clear that approximate matching is obtained from that study as from this. Note also that the deviation from matching is in the direction below the negative diagonal. This results because the programmed omission values were invariably somewhat greater than the obtained values.

The data in Figure 11 might as readily be viewed from the perspective of obtained conditional probability as from obtained omission probability. In the data from the present study, however, these conditional probabilities remain very low throughout most of the range studied, and so a functional relationship is not available However, the question of

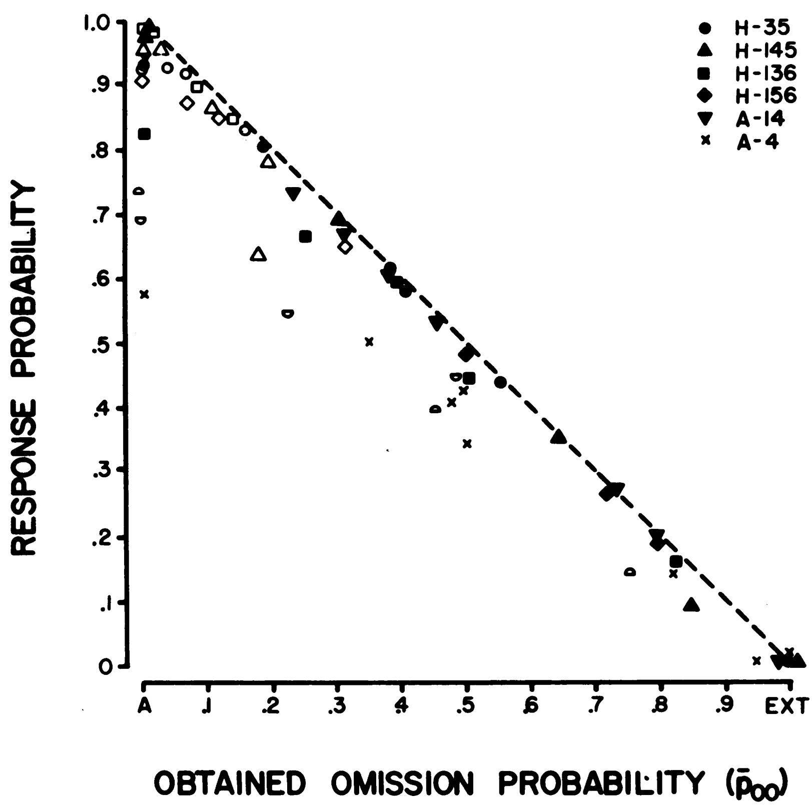

Fig. 11. Response probability as a function of the obtained omission probability ($\bar{p}_{00}$). Open symbols are data from Neffinger and Gibbon (1975); filled symbols from the present study are explained in the legend. The dashed negative diagonal is the matching relation.

which variable is primary in controlling responding here remains moot. The analysis suggested emphasizes the conditional probability of shock given no response, as an all-or-none variable controlling avoidance responding. The approximate matching we observe then results because subjects' response levels fluctuate in a manner which preserves this low conditional probability of shock given no response. One might as easily argue that the primary relationship is the approximate matching in Figure 11. On this view, subjects in the Neffinger and Gibbon study showed little change with changing conditional probability of shock given no response because they approximately matched their response levels to $\bar{p}_{00}$ and ceased responding only when the experimental manipulation of conditional probability of shock given no response was set so low that obtained probabilities increased substantially. Put another way, one might argue that subjects match obtained omission probabilities at a variety of experimentally set conditional probabilities, and the approximate constancy of response probability is reflected in the fact that the open points in Figure 11 all generally fall in the upper left corner of the figure when obtained omission probability is low.

The contrast between the matching view and the all-or-none threshold view parallels in some ways the molar and molecular contrast theories of appetitive control. A molar view in the present case would argue that the matching relation in Figure 11 is primary, while the molecular view would argue that responding fluctuates in such a way as to result in an obtained conditional probability of shock given no response which closely brackets some low "threshold" value.

The manner in which subjects contact the relevant shock frequency variable—whether it be the conditional or absolute probability—evidently requires a substantial number of trials. The data on trial sequences (Figure 7) and on postshock latency distributions (Figures 6 and 10) provide evidence that response probability does not change much on a trial-by-trial basis. This path independence means that the mechanism for adjusting responding to programmed omission probabilities must operate on a fairly large sample of subjects' recent past history.

Punishment

The absolute probability of response-plus-shock produced very different effects from the omission variable. While there was some decrement of response levels with increasing p_{11}, nothing like a linear dependence was observed, and subjects were highly idiosyncratic in the way in which they responded to punishment. Some subjects showed little sensitivity to increasing punishment frequency, and others were very sensitive to punishment. Similar idiosyncracies were observed by Neffinger and Gibbon (1975) under the introduction of postresponse shocks in their punishment subjects. In that experiment, a dichotomous population of subjects was observed, with some showing little reduction in responding with increasing shock probability given a response and others showing eventual cessation of responding with sufficiently frequent shock on response trials. Order effects similar to those observed here were also obtained.

Response probability under the 100% noncontingent shock condition for animals in the Punishment Group, and for the omission subjects studied later under this condition, was distributed over a fairly broad range. However, responding under the traditional avoidance contingency was maintained at a substantial level for all subjects—a level too high to argue for support of this behavior by shock density alone. Thus, the subjects in this study showed shared control between the shock omission contingency and overall shock density when studied under noncontingent shock. The idiosyncratic reaction to shock on response trials is to be contrasted with the very similar behavior of all subjects under control of the shock omission variable, both when that is programmed in the absolute manner studied here or in the conditional manner studied by Neffinger and Gibbon (1975).

Latency data from the present experiment indicate relatively permanent temporal control, independent of overall response levels. Neffinger and Gibbon showed that, even during extinction when responding was declining, the latency distributions showed the same level of temporal control observed during avoidance training. Hineline and Herrnstein (1970) (see also Hineline, 1977) report a similar independence of timing and response probability. In the present case, responding late in the trial cycle had the same character independently of response levels or the manner in which those response levels were generated—that is, independently of the contingency programmed. Timing did show some effects of shock on the previous trial with shocked trials frequently resulting in lowered response probability early in the following trial. This is in contrast to the typical finding of "bursting" after shock in free-operant avoidance (Boren, 1961). Neffinger and Gibbon also found a high frequency of short latency responding in the subjects that maintained responding under noncontingent shock. Our observations parallel theirs for late trial responding, but the short latency pattern was not reliably observed here.

Neffinger and Gibbon utilized a screening procedure which did not involve shaping. Their subjects were required to pass a criterion under the traditional avoidance contingency without pretraining. The result was a high attrition rate (70%) and the authors argued that the screening selected for subjects which showed either a tendency to respond to shocks or a high sensitivity to the shock omission contingency. In the present study, the shaping procedure resulted in an attrition rate of about 30%, and these subjects showed a range of performance under noncontingent shock which was broader than the bimodal population observed by Neffinger and Gibbon.

In summary, idiosyncratic patterns of responding develop when responding which has previously been effective in eliminating shock is subsequently made less effective by shock following response trials. But when responding maintains its shock omission consequence, and not responding is less frequently shocked, performance looks very much alike on either the conditional probability partial schedule or the absolute probability partial schedule.

We would like to suggest that similarity under omission and differences under punishment may be related to similarity in conditioned behavior at different levels of aversive motivation and to variance in aggressive or agonistic behavior following ineffective ("frustrative") response trials. When subjects are faced with less frequent shock on nonresponse trials, one might regard the motivational level for responding as reduced. Gibbon's (1972) quantitative account argues that this motivational discrimination operates in an essentially similar manner for all subjects. The present data are consonant with this view.

When response trials occasionally result in shock, a new element is introduced. We believe this change is related to the degree to which a given individual's response to aversive stimulation following an attempt to avoid it varies—probably in an aggressive manner. Subjects trained on the avoidance contingency have learned that they may reduce noxious stimulation by responding. When these responses are no longer effective and aversive stimulation is delivered anyway, idiosyncracies in the persistence of responding are evident. Some portion of the lever responding observed in this situation is under the control of shock delivery alone. Whether responding here should be characterized as "aggressive" or represents persistence of the previously learned, but now less effective, avoidance behavior is moot. Aggressive responding in the casual observational sense of that word (e.g., biting or striking at the lever) was not observed in our animals after long training. Also, the degree of temporal control under punishment remained comparable to temporal control under the avoidance contingency. Under noncontingent shock, this meant that well-timed responses occurred in close temporal proximity to shock at the end of the trial.

REFERENCES

Boren, J. J. Isolation of post-shock responding in a free-operant avoidance procedure. *Psychological Reports,* 1961, **9**, 265-266.

Boren, J. J., & Sidman, M. Maintenance of avoidance behavior with intermittent shocks. *Canadian Journal of Psychology,* 1957, **11**, 185-192.

Catania, A. C. Elicitation, reinforcement and stimulus control. In R. Glazer (Ed.), *The nature of reinforcement.* New York: Academic Press, 1971.

Church, R. M. Response suppression. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior.* New York: Appleton-Century-Crofts, 1969.

Gibbon, J. Timing and discrimination of shock density in avoidance. *Psychological Review,* 1972, **79**, 68-92.

Gibbon, J., Berryman, R., & Thompson, R. L. Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior,* 1974, **21**, 585-605.

Hineline, P. N. Negative reinforcement and avoidance. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior.* Englewood Cliffs, N.J.: Prentice-Hall, 1977.

Hineline, P. N., & Herrnstein, R. J. Timing in free-operant and discrete trial avoidance. *Journal of the Experimental Analysis of Behavior,* 1970, **13,** 113-126.

Neffinger, G. G., & Gibbon, J. Partial avoidance contingencies. *Journal of the Experimental Analysis of Behavior,* 1975, **23,** 437-450.

Rachlin, H., & Burkhard, B. The temporal triangle: Response substitution in instrumental conditioning. *Psychological Review,* 1978, **85,** 22-47.

Seligman, M. E. P., Maier, S. F., & Solomon, R. L. Unpredictable and uncontrollable aversive events. In F. Robert Bush (Ed.), *Aversive conditioning and learning.* New York & London: Academic Press, 1971.

Received March 18, 1978
Final acceptance November 20, 1978